Will Page

A Case for Completion

“The debate on how to pay artists for their streaming is stuck. Time for a new approach.”

Will Page, Author of Pivot and former Chief Economist of Spotify

“Thanks to TikTok, wedding bands only play two minutes of a song, as part of a medley, because they don’t think people have the attention span for full songs anymore.”

Fred Goldring, Entrepreneur and Entertainment Lawyer

TLDR:

Completion means songs that are streamed in their entirety are worth more than songs that are skipped before the end

Completion would complement the current royalty accounting system and could easily work alongside user- or artist-centric

Modeling shows that long songs are completed as often as short ones, providing assurance to stakeholders

A fourth way?

Everyone agrees: the music streaming economic model needs to change. Yet no alternative to the status quo, which has remained unchanged for two decades, has captured the consensus. I wish to unlock this gridlock.

Advocates of the predominant ‘pro-rata’ system are at odds with proponents of its primary two alternatives, ‘user-centric’ and ‘artist-centric.’ To restore alignment, let’s first reset the objective: if we seek to make fair, fairer, how do we depart from the pro-rata model without causing harm?

Advocates of change are quick to tout the benefits of their preferred approach, but all have downsides. Migrating to either a user- or artist-centric system would entail significant transitional costs. A typical streaming service provider has over 900 contractual agreements, each tied to complex royalty accounting systems for the two main copyrights – the sound recording and composition – as well as the sub-rights for reproduction, performing, communicating, and making available.

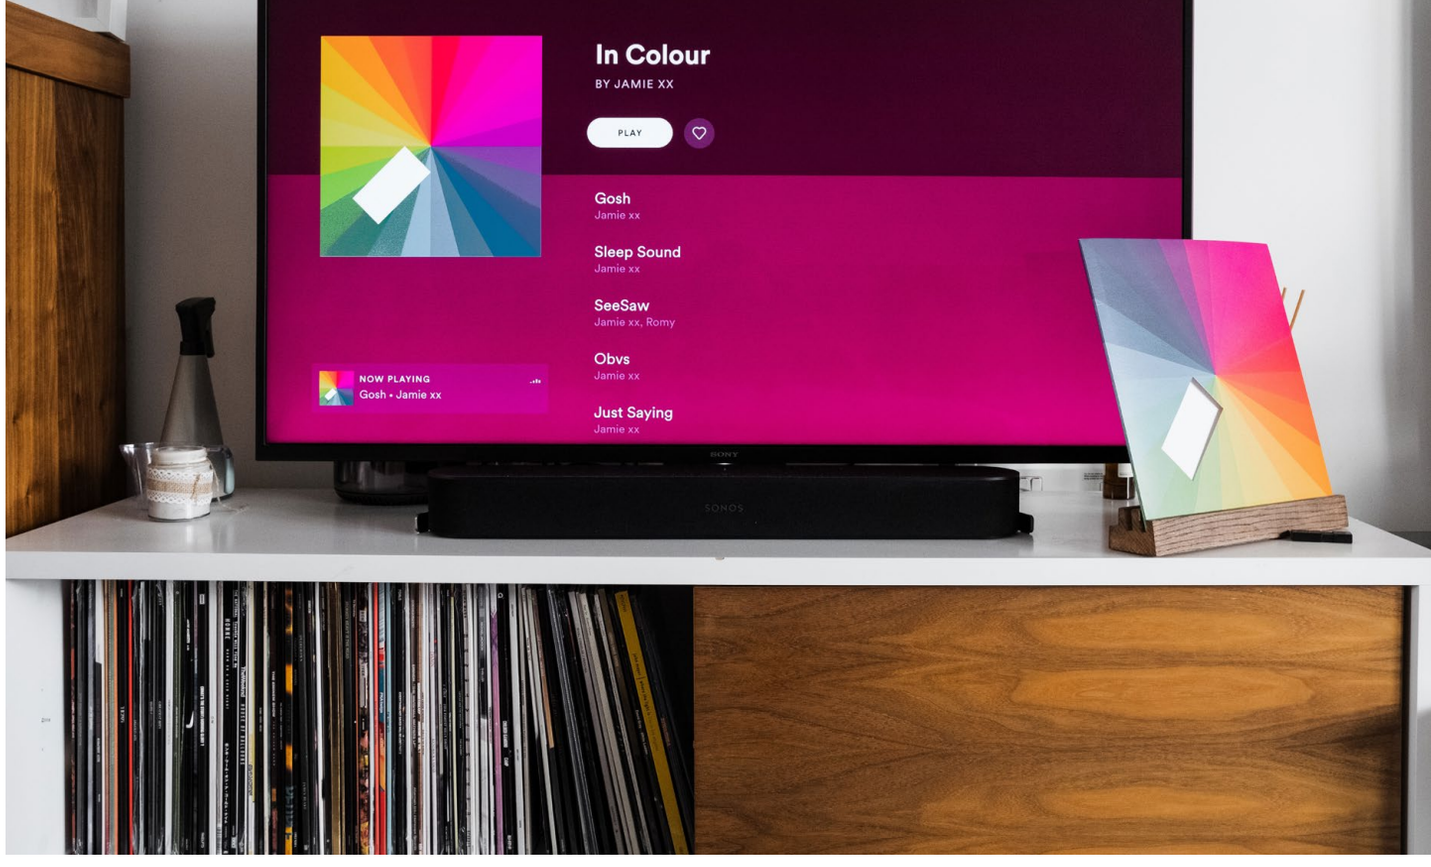

The best case scenario for successfully renegotiating all of this would be lengthy and expensive; worst case would be audit hell.

But there is a way forward: one that not only allows us to move past this impasse, but also allows for iteration once it’s in place. This fourth way is Completion: put simply, songs that are streamed in their entirety should be better compensated, and songs that are skipped before the end should be valued less.

The best argument for this approach is not that compensating completion overcomes the pitfalls of pro-rata, nor that it combats streaming fraud – both of which it does. What makes this proposal sing is that it is practical: the industry can implement it without any contractual or accounting headaches.

This fourth way is Completion: put simply, songs that are streamed in their entirety should be better compensated, and songs that are skipped before the end should be valued less.

Getting proactive with pro-rata

The pro-rata model has served us well for the past decade, as streaming revenues swelled from less than $1 billion in 2012 towards $20 billion by 2022. But agitation is rising. Headlines like *Your Spotify and Apple Music subscriptions pay artists you never listen to* underscore the problem of pooling everyone's subscription revenues together: lighter users subsidise the heavier. Then there's the thorny issue of all songs being worth the same – whether it be sleep music or a songwriting masterpiece.

Five years ago, I co-authored with David Safir two academic papers – 'Money In, Money Out' and 'User-Centric Revisited' – as a prelude to this debate. Our message: learn from century-old collective management organisations (CMOs), which have never assigned the same value to each piece of music. Instead, to make a fair division fairer, CMOs apply rules and qualifications like 'time of day weighting' – which would address the aforementioned sleep music dilemma. Point being, treating music differently is far from unprecedented.

In 2018, Spotify launched Discovery Mode, wherein the streaming platform and (willing) record label partner opted into an agreement to reduce the per-stream payment for a recording, in return for an increased volume of streaming. This 'steering' echoed webcaster Pandora and indie rights aggregator Merlin's 2014 agreement to lower payouts in exchange for a higher volume of plays. Critics may cite 'payola,' but this isn't legally accurate; agreeing to a reduction in payment is not unlawful, and is in fact democratising – effectively enabling those without marketing budgets to do marketing.

In 2021, during the height of the three-year UK Government inquiry into streaming economics, the independent-label trade body AIM commissioned me and Safir to build out their Artist Growth Model. AIM's proposal was to redistribute earnings from the most-streamed artists and spread them more evenly across the market to help niche and emerging artists earn a living. A succinct political summation would be to tax the rich and distribute to the slightly less rich – arguably a form of progressive taxation.

Soon after, Soundcloud launched Fan Powered Royalties, an adaptation of user-centric payments, where artists who opt in receive payouts based on an allocation of users' subscription fees. Beyond the financial implications, this connected artists more directly to their fans. US rapper Lil Uzi Vert, for example, not only could observe that the top 7% of his fanbase was contributing 70% of his revenues; he could also directly market to those superfans as well. You can't discuss the current phrase du jour, 'fandom,' without factoring in fan typology, and the targeting opportunities it affords. Nor can you harness this under pro-rata, where the law of averages means all streams (from all fans) have to be worth the same.

More recently, Deezer and UMG announced they would implement artist-centric payments, with the stated dual-objective of penalising the 'clutter' that pollutes the platform and rewarding the songs that listeners actively seek. In contrast to the aforementioned progressive taxation scheme, artist-centric taxes the very poor – those with less than 1,000 streams and 500 followers – to give to everyone else. Spotify has since modernised their royalty system, where tracks need to reach 1,000 streams in the previous (rolling) 12 months to generate royalties.

One of the most notable features of the artist-centric proposal is that actively searched ('pull') streams would be worth more than those served by the platform ('push' streams). Arguably, this makes sense based on a presumed utility to the listener. But we shouldn't ignore a lesson taught to us by David Safir at a recent NYLON conference, when he asked us all to step back and ask **'for whom are we defining fairness?' Is it the label, the platform, or the consumer?**

If the consumer opts into a contract for the benefit of lean-back 'push' music, who are we to disagree?[1]

1 Frederik Juul Jensen (Université Sorbonne Paris Nord) offers a robust evaluation all the various models in: Rethinking Royalties – Alternative Payment Systems on Music Streaming Platforms

A time-conscious tour of duration

So if the task of making 'fair division' fairer is stuck in a holding pattern, how do we release it? At the DCMS select committee hearings, I pitched the concept of duration. This has precedent: when the BBC pays for the right to exploit music on radio and television, it pays the rights organisation PPL, which then applies a value-per-second rule to its payout distribution. Bohemian Rhapsody (5 minutes and 55 seconds) will earn double that of, say, You're My Best Friend (2 minutes and 52 seconds).

Look around and you'll see other methods of applying duration. The German collecting society GEMA uses ranges, wherein songs that last 6-12 minutes get paid more than those lasting 0-6 minutes. The Mexican collecting society SACM applies a 'taxi fare' sliding scale: the second minute is worth less than the first. The US Mechanical Licensing Collective should, in theory, apply an overtime bonus to longer streams (see box).

Going into overtime

Did you know that in the United States duration for streaming is already codified by law? Look up the small print hidden in Federal Statute §385.21, and you'll find the words 'overtime adjustment' – a relic from the CD world that should, in theory, be applied to streaming. Songs are adjusted as follows:

(1) **5:01 to 6:00 minutes** – Each Play = **1.2 Plays**
(2) **6:01 to 7:00 minutes** – Each Play = **1.4 Plays**
[...and so on...]
(6) For playing times of greater than **10 minutes**, continue to add **0.2 Plays** for each additional minute or fraction thereof.

From theory to practice: in the age of the CD – when labels passed through the mechanical royalty to publisher – the longer the song, the more the publisher received, and the less the record label retained. In reality, though, this was suppressed through 'controlled composition' clauses.

In streaming, there is no pass-though, and streaming platforms are not required to pay more for longer songs; hence, overtime is not activated. But if it were, it would be demand-reallocation, not creation – shifting money away from shorter tunes to pay for the longer ones.

So, in theory, duration already exists in the U.S., but in practice we don't watch the clock.

A third threshold: capturing completion

Duration is thus neither a novel idea nor as one-dimensional as it first appears. Nonetheless, in applying it to streaming, it's tempting to recall the Scottish tourist office expression: 'if you wanted to get there, you wouldn't start from here'. Calculating duration on a domestic radio station's one-to-many log sheet is easy; on streaming platforms, calculating the length of each song, and the consumption of each song, for every unique user on the platform, in real time, is hard. Damn hard. We need baby steps.

But here's the plot twist: we already measure duration on streaming! Only we measure it as a threshold: listeners must play a song for an uninterrupted 30 seconds for a royalty to be payable.

Skip before 30 seconds and it won't count to the rights holder's payout (or the streaming platform's costs). Arguably there are actually two thresholds in place: first, the legal and technical definition that a song has been played; then second, whether 30 seconds has elapsed. So, let's skip the complexity of comprehensively measuring listening-duration and instead simply introduce a third threshold: completion.

The case for completion – rewarding songs that are listened to in full and penalising those that are skipped – can be achieved at relatively low marginal (and transitional) cost. Implementing this third threshold is no more complicated than the first or second. It's easy to legally define, and therefore removes the risk of auditing disputes. And it not only fulfils the widely held desire to depart from pro-rata without causing harm, but also allows for more aspirational models to be explored.

Of course, any proposed alternative to pro-rata will be met with the same knee-jerk question: "Who are the winners and losers?" This is often followed by the rather naïve retort of, "Well, that wouldn't be possible, because it would create losers." To be clear, pro-rata has winners and losers, too, and they change over time. What's more, you can't do retrospective analysis on who the winners and losers might be under a future model, since when you change the rules, you change the way the game is played.

Change the rules and you change the game

Good news: MusicAlly, a consultancy, lists over 20 studies on user-centric payouts on its helpful website. Bad news: they're all now null and void. Retrospective studies like these that apply user-centric distribution on a pro-rata behavioural basis fail to capture how changing the rules changes the game.

As explained in SoundCloud Rockonomics, we've already seen this dynamic play out in streaming. Why are songs getting shorter and choruses being moved to the front? Because the rules say you don't get paid unless 30 seconds of your song have been played, and not a penny more for anything longer.

Like a tail that wags the dog, so too do the rules of streaming dictate the game of payouts.

To make matters even more confusing, if a new royalty model were to be implemented and neither the creator nor consumer were aware the rules had changed (which is not implausible), then we would have to be careful about inferring anything from the experiment.

Nevertheless, if the case for completion is to be considered by an industry that's already proven reluctant to change (and at times cynical toward proposals), **we should be explicitly clear on its strengths, weaknesses, and potential unintended consequences.** The following is not exhaustive, but turns over the stones that most obviously need to be examined.

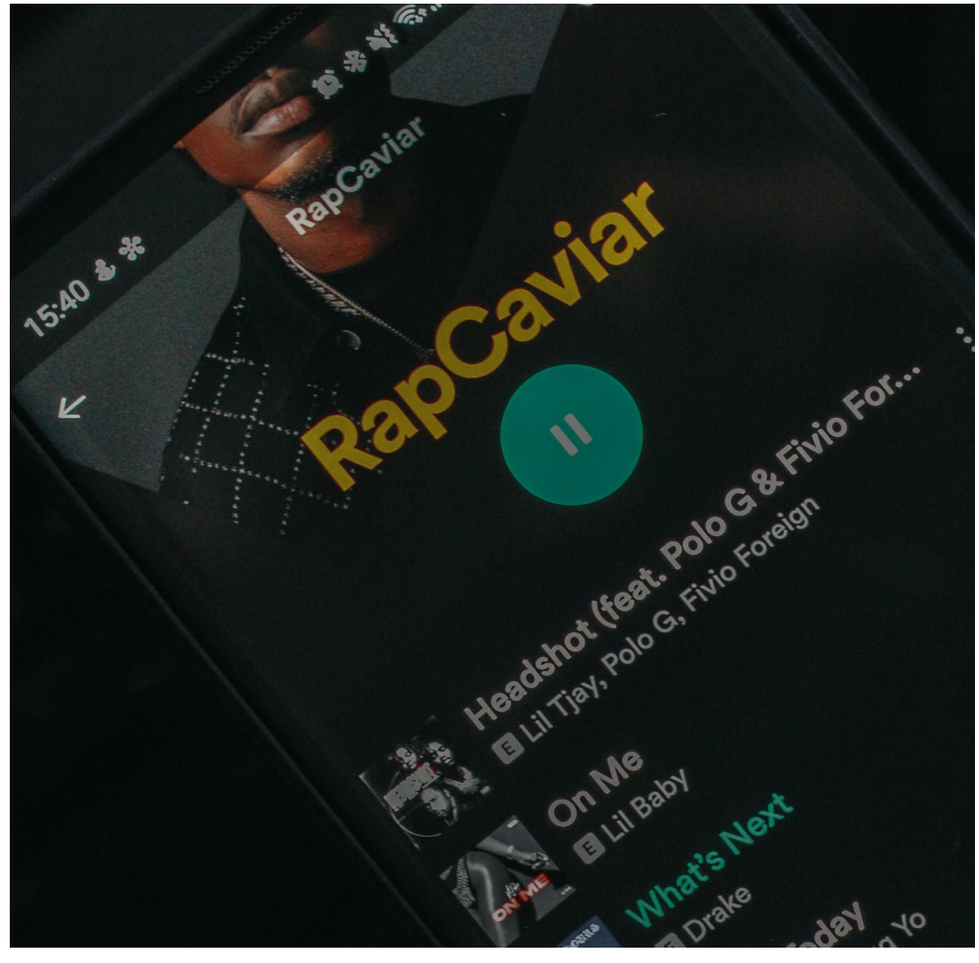

Strengths of Completion

The primary strength of completion is its common sense: if a listener stops listening to a song halfway through, then it generates fewer royalties than a listener who listens to the whole thing. Moreover, the platform can only penalise incompletion by the consumer's actions, so it's opt-in by nature.

Other strengths of completion include:

- Culturally, it rewards greater appreciation of the songwriter's work, and acts as a counter to today's swipe-and-skip mentality toward the arts.
- Economically, it incentivises music to compete for attention that might otherwise be lost to other media merchants (c.f. Netflix views its biggest competitor as sleep).
- Indirectly, it rewards intent, as listeners don't intervene to complete, but do so to skip.

- Legally and technically, it's a threshold that's easy to define, implement, and audit (although endings can be unclear, as I discovered with my prepubescent introduction to Led Zeppelin's 'Thank You' from their seminal second album).
- Fraudsters can no longer run click-farms based on mass-repeats of 31-second plays.

Weaknesses of Completion

The main weakness of completion is that it unintentionally favours lean-back listening, and 'inadvertent dormant listening', like when you've left Spotify running while you take a shower. This isn't bad per se (the consumer has just as much right to define fairness as the creator), nor is it that different from radio – but it is a weakness.

Other *potential* weaknesses include:

- Could punish longer songs, which may have a higher propensity to be skipped, a point we'll counter with evidence later.
- May struggle to measure scrubbing, where listeners shuffle back or forth ten seconds.
- Incentivises songwriters to pen shorter songs, where skipping may be perceived as being less likely – again, a misconception that we'll correct later.
- Stumbles on a rediscovery tradeoff, as you may skip what you already know.
- 'Fraudsters gonna fraud' – and may start releasing short songs (50-70 seconds) that game the three thresholds of (i) play, (ii) thirty uninterrupted seconds, and (iii) completion – although this would now be easier to detect.

Unintended consequences of Completion

The introduction of a third threshold has the unintended consequence of opening the floodgates to more complexity, such as additional thresholds at (say) 90 seconds, or 120 seconds, to determine how much of a song a listener has heard. Simplicity begets complexity, reminiscent of the Sorites Paradox.[2]

Other unintended consequences of this proposal include:

- Creates a numerical anomaly, where the completion of pop songs lasting two minutes is worth more than the incompleteness of a six-minute song, raising questions about the ideal duration of a song.
- Incentivises labels and publishers to focus on promoting shorter songs – and it could mean playlists fill with shorter songs, which may punish older content and certain genres.
- Completion doesn't negate user/artist-centric; indeed, this proposal could sit alongside both of these options (and many others).
- Incentivises how songwriters choose how to conclude their songs, and may result in fewer Elvis Costello long fade-outs (which is not necessarily a bad thing!).
- Doesn't remove the signal-to-noise ratio dilemma, as we would still have no true measure of the appreciation of music (e.g. signals like 'follow' are meaningless, and 'likes' lack meaning); nor would we yet fully capture the option value of streaming: that it's not what you do with the service, but what you could do that often matters more.

2 At the time of writing, Netflix has just informed your author that his monthly fee is going to be raised to £17.99. I don't intend to churn but would note that it's now more than doubled and I use the service even less.

Modelling the case for completion

To model completion, and get a sense of how it would play out in practice, we use here three months' worth of data on the top 15,000 songs on UK streaming services as measured by the Entertainment Retailers Association (ERA) – enough to capture 52% of all streams in the UK during this period. These data include each song's label, distributor, and corporate group, along with genre and release date. Additionally, since these data include ISRC codes, blending them with the Echo Nest and MusicBrainz data allows for the extraction of song length for each song.

With modelling, it's important to keep in mind that (a) it is a retrospective exercise, and (b) in this case, it's just a one-country data set. The results will differ for specific streaming services both within the UK and across other markets, as well as on different tiers (free and paid, as well as audio and video), and during different seasons.

Mapping Song Length to Completion

For this model, we estimate the likelihood of completion for a given song length, then use those estimates to assess the per-stream payout results of punishing incompletion. **As a model 'driver,' song length makes intuitive sense, as it is the one thing the creator controls** – though for thoroughness, we also consider some factors outside of the creator's control, namely source of streams (push or pull) and genre.

To estimate the likelihood of completion given a song's length, we first use data from the Echo Nest on completion rates by duration – that is, the likelihood that a listener completes a song within a given listening time frame. These probabilities are then used as a proxy for completion rates by song length.

The results are shown in one-minute intervals, up to ten minutes long, with one additional bucket for durations/song lengths beyond ten minutes. Whilst mapping duration onto song length isn't perfect, it's intuitively reasonable: if a song-listening session that lasts between 4 and 5 minutes is 95% likely to result in a completed song, it is sensible to use the same probability for the likelihood that a song lasting between 4 and 5 minutes will be completed.

There are two obvious outliers in these data: the anomalously low probabilities of completion for duration/song lengths lasting under one minute and between 1-2 minutes. In addition to the relatively low number of songs of these lengths, these low probabilities capture the prominence of skipping in the case of the former, and interludes and interstitials in the latter.

At all longer intervals, completion rates hover around ninety percent, with no evidence of decay as song length climbs, suggesting some **self-selection is at play – those with shorter attention spans selecting shorter songs, and those with longer choosing lengthier.**

COMPLETION RATES BY DURATION RANGE IN MINUTES

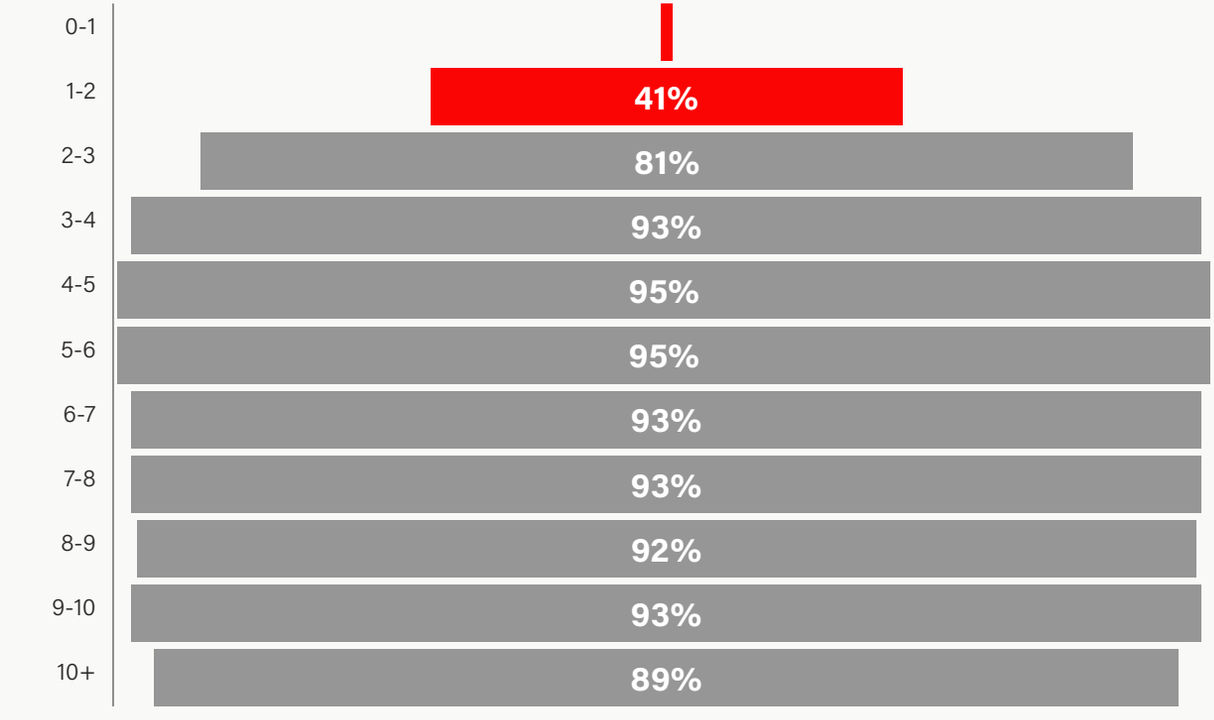

Source: Echonest

The Completion Model

The best way to think about the mechanics of any model is to work through its inputs, variables, and outputs.

Inputs

First, we need to feed our model an input – in this case, a fixed pot of cash that streaming services generate for record labels (publishers are omitted from this analysis). To make this input as realistic as possible, we use figures broadly in line with income reported by the UK trade bodies ERA and BPI. We assume British consumers (and their associated ad revenues) generate £110m in streaming revenues per month (inclusive of VAT). Strip out VAT at 20%, and you're left with a licence base of around £90m, of which labels capture an effective revenue share (ERS) of 55%, resulting in our input of £50m in payments for record labels.

Variables

Next is the component that the streaming platform can control, 'completion premium,' which dictates how much a completed song is rewarded relative to an uncompleted song. We've used effective revenue share (ERS) as the variable defining the completion premium. In this model, the only adjustment needed is the ERS for incompletion, which then implies the completion premium. Notably, since the total share of streams that are incomplete is only about 10 percent, the adjustment needs to be relatively significant in order to provide a material reward to completed songs.

Outputs

By applying an incompletion punishment of 15 percentage points, lowering the ERS from 55% to 40%, and transferring that lost revenue to the completion premium, we find completed streams receive an ERS of 58%. Compared to the status quo per-stream revenue of £0.0048, completed songs get a 3% uplift to £0.0050 per stream, while incomplete streams get a 27% reduction to £0.0035.

The model – and its implementation – allows for a range of completion premia and incompletion punishments. Taken to its extreme, treating incompletion like a below-30 second stream and declaring these 'non-royalty bearing,' the completed per-stream revenue would rise to £0.0054 (an uplift of 12.1%). What you can't adjust is the input: the £50m that's paid over from streaming services to labels each month remains fixed. This is a demand allocation model, not demand creation.

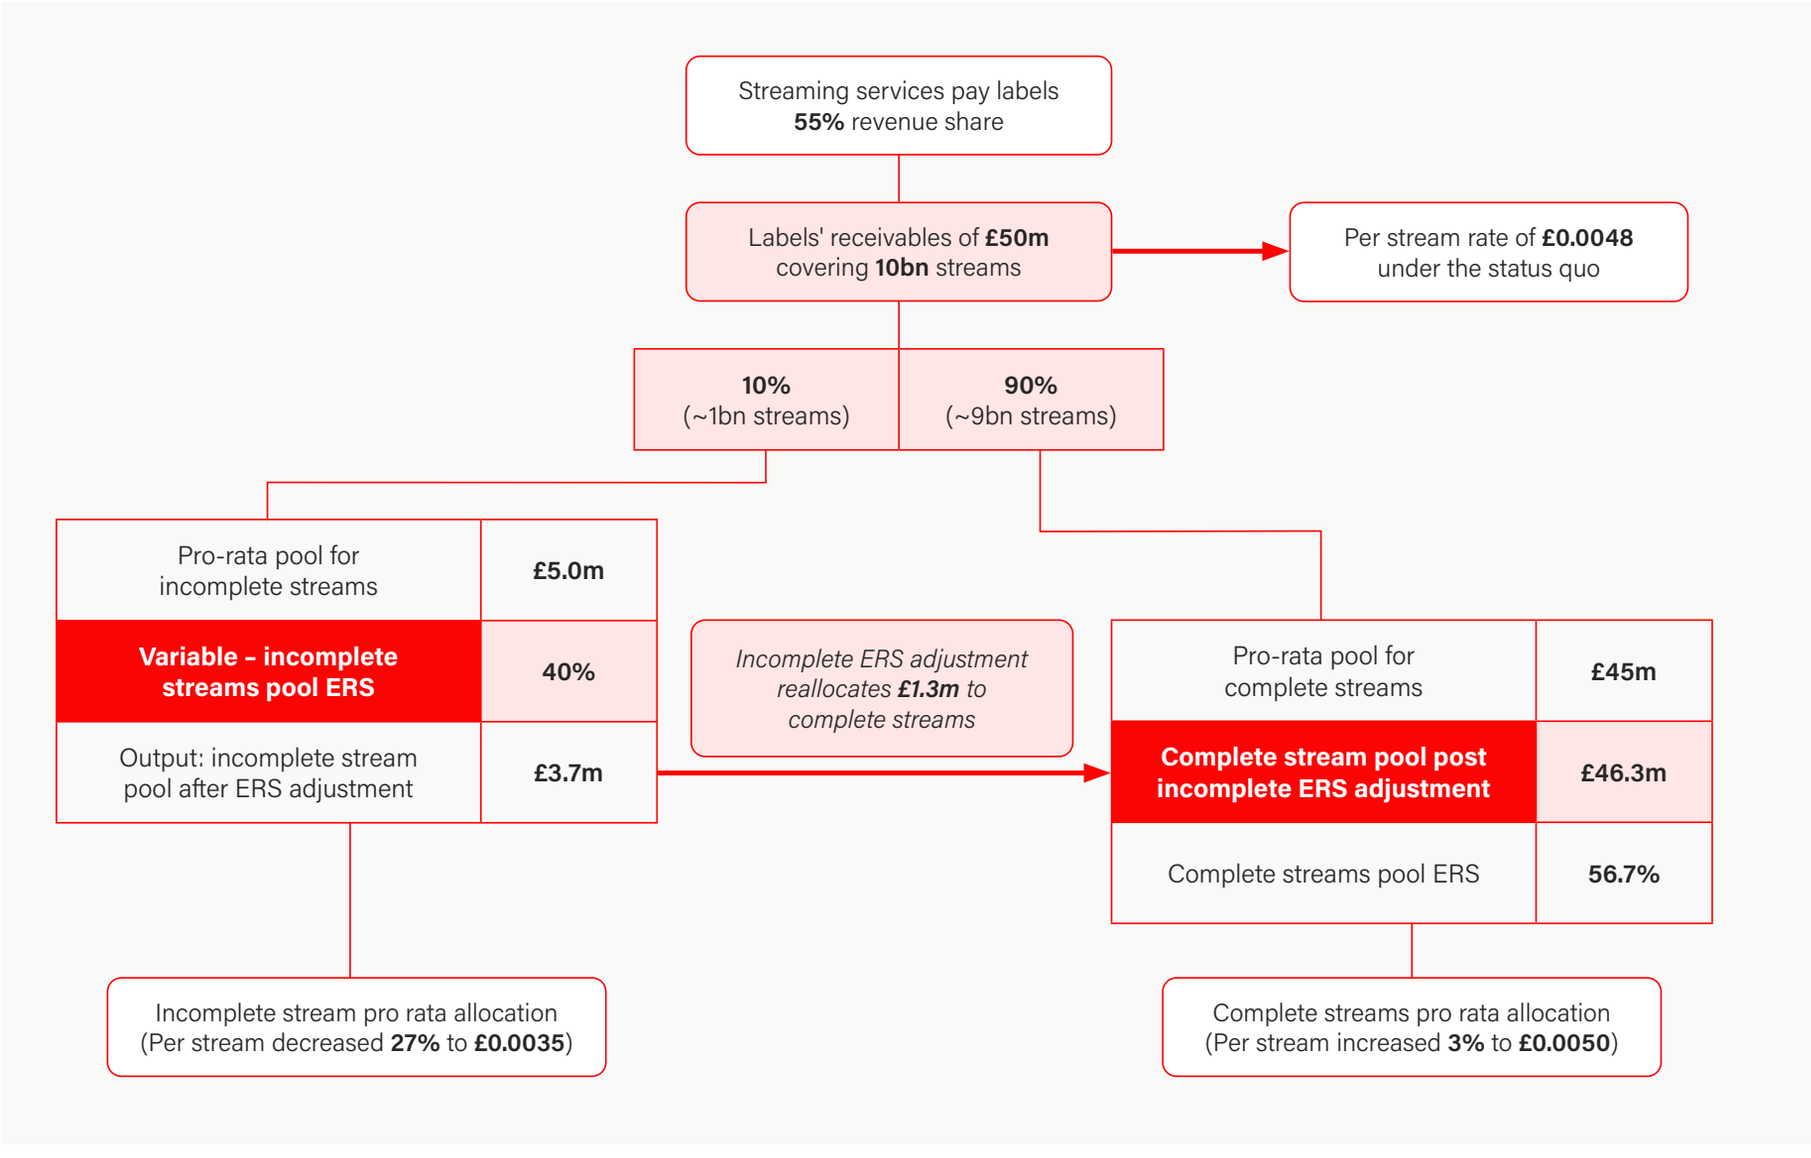
Streaming services pay labels **55%** revenue share
Labels' receivables of **£50m** covering **10bn** streams
Per stream rate of **£0.0048** under the status quo
10% (~1bn streams)
90% (~9bn streams)
Pro-rata pool for incomplete streams
£5.0m
Variable – incomplete streams pool ERS
40%
Output: incomplete stream pool after ERS adjustment
£3.7m
*Incomplete ERS adjustment reallocates **£1.3m** to complete streams*
Pro-rata pool for complete streams
£45m
Complete stream pool post incomplete ERS adjustment
£46.3m
Complete streams pool ERS
56.7%
Incomplete stream pro rata allocation (Per stream decreased **27%** to **£0.0035**)
Complete streams pro rata allocation (Per stream increased **3%** to **£0.0050**)

Going to Further Lengths

'Whataboutism' always prevails when you complete a model and pause to think about what might be missing. One factor is genre, which is defined after a song is released and its definition can vary by market. The Echo Nest points to jazz and classical as having some of the highest completion rates, whereas pop has some of the lowest. A drawback to assessing completion by genre is that the Echo Nest data doesn't consider the popularity of songs within their genres.

Point being, popular songs that climb to the top of the charts are more likely to be completed simply because they are popular – that's how music works.

Another 'whataboutism' that'll be high on readers' minds, but not featured in our model, is whether the stream is a 'push' or a 'pull.' The ERA data doesn't include this, nor can the streaming platform or creator control for it. The Echo Nest can offer some context here to satisfy our curiosity, however, bucketing streams into five categories, as shown below, with red indicating lean-back and grey capturing lean-forward. Completion rates are, unsurprisingly, highest for the 'lean-back' radio feature (82%), and lowest for 'lean-forward' non-playlist listening (73%).

COMPLETION RATE BY SOURCE OF STREAMS

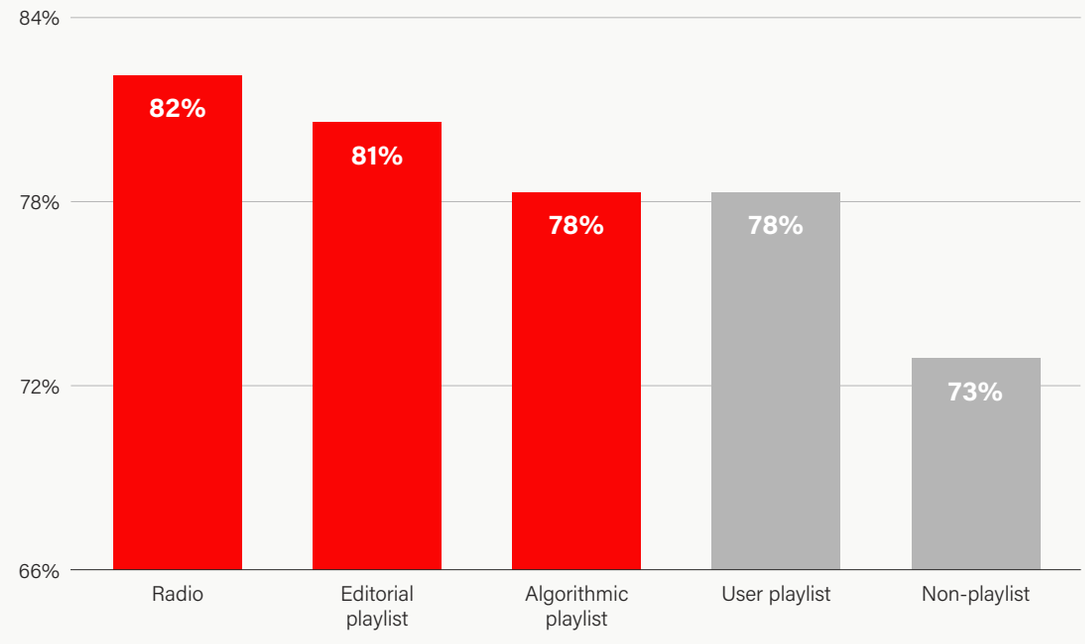

Source: Echonest, Spotify

Concluding Completion

Completion is the best way for our industry to evolve from pro-rata without collapsing current royalty systems. It doesn't preclude other proposals, which can follow, so it is best thought of as a means to a fairer end, not an end in itself.

Importantly, completion is relatively easy to implement. All it asks for is the introduction of a third threshold to accompany the already extant first (defining a play) and second (ensuring thirty uninterrupted seconds have been streamed). This can be done with little marginal (or transitional) costs.

Completion penalises intent, as listeners don't intervene to complete a song, but do so to skip. We can safely attribute this intentional act to a lower valuation by the listener of the music being skipped. Rewarding songs according to the listener's valuation of them is a workable definition of fair. There are spillover benefits to rewarding completion, too, especially in tackling stream fraud (31 second stream-farms would be killed overnight) and refocusing an industry toward competing for scarce attention – which, in an age where we struggle to screen out distractions, might otherwise be lost.

Finally, completion doesn't compete with other potential royalty changes – it complements them.

Those who wish to continue exploring user-centric may recognise a benefit that the 'my money, my music' approach is being fine-tuned here to reward 'my completed music with more of my money.' Artist-centric advocates can take heart, too, as there's a clear correlation between actively searched streams and completion – perhaps worthy of an 'octuple boost'. Finally, those who simply want to make inroads into fighting fraud and removing bad actors from streaming can monitor and evaluate the incremental costs and benefits of introducing this third threshold before considering the need for further measures.

In recognition of an important lesson once taught to me by the historian Simon Schama, who rose to fame writing ‘A History of Britain,’ I’d like to draw attention to the choice of words in the title of this essay: “A” case for completion, not “The”. I don’t claim to own this idea, and welcome competing alternatives.

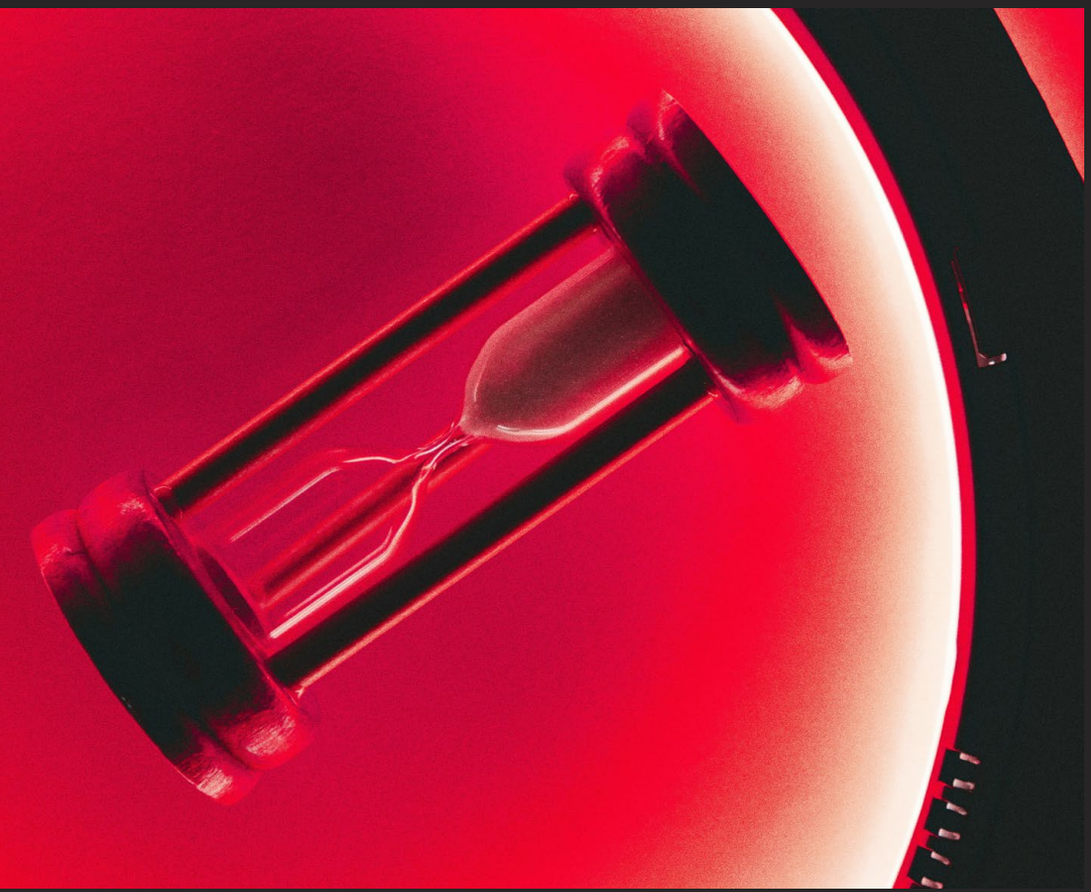

What this work brings to the debate is an evidence-based approach to fairness that is both holistic in its approach and easy for creators and executives to understand.

Too often, fairness trades off with efficiency – but on this occasion I believe we’ve found a welcome balance.

Acknowledgements

The author would like to acknowledge the wordsmith extraordinaire Sam Blake for copy-editing the case for completion and making it readable all the way to the very end. Alice Clarke's outstanding design skills practice what the proposal preaches – as if you're reading this, you've completed the author's work. He also needs to appreciate the exceptional modelling skills of Nicholas Lightle, Chris Tynan and the music 'alchemist' Glenn McDonald (formerly of the Echo Nest) which made this work possible, and the cooperation of the Entertainment Retailers Association for providing the data. The author would also like to acknowledge the influence of David Safir, who has been both a mentor and an inspiration throughout this never-ending debate about fairness. The author is grateful to: Thomas Breslin (ReRight); Benedict Curran (Decca Records); Morgan Hayduk (BeatDapp); Frederik Juul Jensen (Université Sorbonne Paris Nord); Ben McEwen (ICE Services); Shannon Nitroy (Spotify) and Paul Sanders (State 51). Finally, a special thanks to Fred Goldring for inspiring this work with his 'wedding planner' observation.